

REMARKS

I. Introduction

In response to the Office Action dated March 11, 2004, no claims have been cancelled, amended or added. Claims 1-36 remain in the application. Re-examination and re-consideration of the application is requested.

II. Prior Art Rejections

A. The Office Action Rejections

In paragraph (5) of the Office Action, claims 1, 3-10, 12, 13, 15-22, 24, 25, 27-34, and 36 were rejected under 35 U.S.C. §103(a) as being unpatentable over Viniotis et al., "Linear programming . . . Queueing systems," IEEE, 1988 (Viniotis) in view of Schneider et al., "Stochastic Production scheduling . . . demand forecasts," IEEE, 1998 (Schneider). In paragraph (6) of the Office Action, claims 2, 14, and 26 were rejected under 35 U.S.C. §103(a) as being unpatentable over Viniotis in view of Schneider and further in view of Dangat et al., U.S. Patent No. 5,971,585 (Dangat). In paragraph (7) of the Office Action, claims 11, 23, and 35 were rejected under 35 U.S.C. §103(a) as being unpatentable over Viniotis in view of Schneider and further in view of Hedlund et al., "Optimal control of hybrid systems," IEEE, 1999 (Hedlund).

Applicant's attorney respectfully traverses these rejections.

B. The Applicant's Invention

Independent claims 1, 13 and 25 are generally directed to a method for solving, in a computer, stochastic control problems of linear systems in high dimensions. Claim 1 is representative, and comprises:

(a) modeling, in the computer, a structured Markov Decision Process (MDP), wherein a state space for the MDP is a polyhedron in a Euclidean space and one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state; and

(b) building, in the computer, one or more approximations from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming.

C. The Viniotis Reference

Vinotis describes linear programming as a technique for optimization of queuing systems. For a significant number of queuing models, that appear in diverse, seemingly unrelated application areas, such as routing, resource allocation and flow control, the optimal policy exhibits a certain "switching-curve" structure. In this paper, we formulate the optimal control problem of such models in a unified way, by using abstract Linear Programming. Using well-known facts from sensitivity analysis of Linear Programs, we show how certain properties of the optimal policy can be easily derived, even in cases where Dynamic Programming (DP) and Stochastic Dominance (SD) arguments fail. A structural property of the optimal value function of the Linear Program, namely piecewise linearity, is exploited to derive properties of the optimal cost function. We also consider additional problems in the realm of queuing system control in which DP or SD approaches are not applicable but Linear Programming may provide useful results.

D. The Schneider Reference

Schneider describes stochastic production scheduling to meet demand. Production scheduling, the problem of sequentially configuring a factory to meet forecasted demands, is a critical problem throughout the manufacturing industry. The requirements of maintaining product inventories in the face of unpredictable demand and stochastic factory output make the problem difficult. Existing approaches commonly fall into one of two groups: either demand forecasts are discarded and linearizing assumptions are made so methods based on optimal control can be applied, or AI search methods are used to tackle the large search spaces and the ability to handle stochasticity optimally is sacrificed. This paper describes a Markov Decision Process (MDP) formulation of production scheduling which captures stochasticity, while retaining the ability to construct a schedule to meet demand forecasts. The solution to this MDP is a value function, specific to the current demand forecasts, which can be used to generate optimal scheduling decisions online. The paper then describes an industrial application and a reinforcement learning method for generating an approximate value function in this domain. The results demonstrate that in both deterministic and noisy scenarios, value function approximation is an effective technique.

E. The Dangat Reference

Dangat describes a computer implemented decision support tool serves as a solver to generate a best can do (BCD) match between existing assets and demands across multiple manufacturing facilities within boundaries established by manufacturing specifications and process

flows and business policies to determine which demands can be met in what time frame by microelectronics (wafer to card) or related (for example disk drives) manufacturing and establishes a set of actions or guidelines for manufacturing to incorporate into their manufacturing execution system to insure the delivery commitments are met in a timely fashion. The BCD tool has six major components, a material resource planning explode or "backwards" component, an optional STARTS evaluator component, an optional due date for receipts evaluator, an optional capacity available versus needed component, an implode "forward" or feasible plan component, and a post processing algorithm.

F. The Hedlund Reference

Hedlund describes optimal control of hybrid systems. This paper presents a method for optimal control of hybrid systems. An inequality of Bellman type is considered and every solution to this inequality gives a lower bound on the optimal value function. A discretization of this "hybrid Bellman inequality" leads to a convex optimization problem in terms of finite-dimensional linear programming. From the solution of the discretized problem, a value function that pre-serves the lower bound property can be constructed. An approximation of the optimal feedback control law is given and tried on some examples.

G. Applicant's Claims Are Patentable Over The References

Applicant's claims are patentable over the references because they recite a novel and nonobvious combination of elements. Specifically, Applicant's claims are patentable because they recite a novel and non-obvious combination of "displaying," "selecting," "mapping" and "invoking" elements. Neither of the references, taken individually or in any combination, teaches or suggests this sequence of steps.

The Office Action states the following:

5. Claims 1, 3-10, 12, 13, 15-22, 24, 25, 27-34 and 36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Viniotis et al. (VI) ("Linear programming ... Queueing systems", IEEE, 1988) in view of Schneider et al. (SC) ("Stochastic Production scheduling ... demand forecasts", IEEE, 1998).

5.1 VI teaches Linear programming as a technique for optimization of queuing systems. Specifically, as per Claim 13, VI teaches solving stochastic control

problems of linear systems in high dimensions (Page 652, CL1, Para 1; Page 653, CL2, Para 3); comprising:

modeling a structured Markov Decision Process (MDP) (Page 652, CL1, Para 4; Page 652, CL2 Para 6), wherein a state space for the MDP is a polyhedron in a Euclidean space (Page 654, CL2, Lemma 2);

one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state (Page 653, CL1, Para 1 and Para 2; Page 652, CL2, Para 7); and

building a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming (Page 653, CL1, Para 9 to Page 654, CL1, Para 4; Page 652, CL2, Para 8).

VI does not expressly teach a computerized apparatus for solving stochastic control problems of linear systems in high dimensions comprising a computer. SC teaches a computerized apparatus for solving stochastic control problems of linear systems in high dimensions comprising a computer (Page 2726, CL1, Para 3 and 4), as that allows the solution of stochastic control problems of linear systems in high dimensions run faster and allows the user to generate the results with varying data (Page 2726, CL1, Para 3). It would have been obvious to one of ordinary skill in the art at the time of Applicant's invention to combine the method of VI with the apparatus of SC that included a computerized apparatus for solving stochastic control problems of linear systems in high dimensions comprising a computer type, as that would allow the solution of stochastic control problems of linear systems in high dimensions run faster and allow the user to generate the results with varying data.

VI does not expressly teach logic performed by the computer, for modeling a structured Markov Decision Process (MDP). SC teaches logic performed by the computer, for modeling a structured Markov Decision Process (MDP) (Page 2726, CL1, Para 3 and 4), as that allows the solution of stochastic control problems of linear systems in high dimensions run faster and allows the user to generate the results with varying data (Page 2726, CL1, Para 3). It would have been obvious to one of ordinary skill in the art at the time of Applicant's invention to combine the method of VI with the apparatus of SC that included logic performed by the computer, for modeling a structured Markov Decision Process (MDP), as that would allow the solution of stochastic control problems of linear systems in high dimensions run faster and allow the user to generate the results with varying data.

VI does not expressly teach logic performed by the computer, for building a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming. SC teaches logic performed by the computer, for building a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming (Page 2726, CL1, Para 3 and 4), as that allows the solution of stochastic control problems of linear systems in high dimensions run faster and allows the user to generate the results with varying data (Page 2726, CL1, Para 3). It would have been obvious to one of ordinary skill in the art at the time of Applicant's invention to combine the method of VI with the apparatus of SC that included logic performed by the computer, for building a value function for the state using representations that facilitate the computation of approximately optimal

actions at any given state by linear programming, as that would allow the solution of stochastic control problems of linear systems in high dimensions run faster and allow the user to generate the results with varying data.

VI does not expressly teach logic performed by the computer, for building one or more approximations from above and from below to a value function for the state using representations. SC teaches logic performed by the computer, for building one or more approximations from above and from below to a value function for the state using representations (Page 2722, CL1, Para 2; Page 2724, CL2, Para 6), as value function approximation is an effective technique for both deterministic and noisy scenarios (Page 2722, CL1, Para 2); and approximation allows solving large scale MDPs (Page 2722, CL2, Para 2). It would have been obvious to one of ordinary skill in the art at the time of Applicant's invention to combine the method of VI with the apparatus of SC that included logic performed by the computer, for building one or more approximations from above and from below to a value function for the state using representations, as value function approximation would be an effective technique for both deterministic and noisy scenarios and approximation allows solving large scale MDPs.

Applicant's attorney disagrees. Neither reference, taken individually or in combination, discloses the specific combination of elements set forth in Applicant's independent claims 1, 13 and 25.

For example, the Office Action asserts that Viniotis teaches "a state space for the MDP is a polyhedron in a Euclidean space," at page 654, CL2, Lemma 2. However, at the indicated location, Viniotis merely states the following:

Viniotis: page 654, CL2, Lemma 2

Lemma 2: If A is a totally unimodular matrix, the extreme points of the polyhedron $\{y: Ay \leq b\}$, where the vector b is integer-valued, are vectors with integer components.

However, in Viniotis, A is a constraint matrix, not a state space. Moreover, Viniotis does not refer to a polyhedron in a Euclidean space.

In another example, the Office Action asserts that Viniotis teaches "one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state," at page 653, CL1, Para 1 and Para 2; Page 652, CL2, Para 7. However, at the indicated locations, Viniotis merely states the following:

Viniotis: page 653, CL1, Para 1 and 2

Thus, any linear cost functional that involves the state (e.g., delay), is linear in the controls z_k . Selecting an optimal policy, therefore, reduces to minimizing a linear functional; this minimization is constrained, since the states generated by the policy have to belong to the state space S , a (possibly unbounded) subset of the

nonnegative integers. From the state equation, the constraints are also linear in the control. But minimization of a linear functional over a linear constraint set is the subject of Linear Programming.

There are some points that need attention. In a Linear Program, the control variables are allowed to take values in a continuum, e.g., $[0,1]$ or \mathbb{R}^n . In (an unconstrained) MDP problem, the controls are integer-valued. For example, in resource allocation problems, where there are $N+1$ distinct actions available, $z_k \in \{0,1,\dots,N\}$. Thus when reformulating the problem as a Linear Program, we in fact "enlarge" the solution space. This will not be a problem if existence of integer-valued optimal solutions is shown.

Viniotis: page 652, CL2, Para 7

In the next section we briefly present the technicalities of the formulation of the MDP problem as a linear program; we use the notation developed in [7]. The reader may find the missing details in [7,14].

In reviewing the above, it can be seen that Viniotis teaches only that a linear cost functional that involves the state is linear. However, these portions of Viniotis do not teach or suggest that actions that are feasible in a state of the state space are linearly constrained with respect to the state, in the context where a state space for the MDP is a polyhedron in a Euclidean space.

In another example, the Office Action asserts that Viniotis teaches "building a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming," at page 653, CL1, Para 9 to Page 654, CL1, Para 4; Page 652, CL2, Para 8. Applicant's attorney notes that the recitation of the limitations are incorrect. Moreover, at the indicated locations, Viniotis merely states the following:

Viniotis: page 653, CL1, Para 9 to page 654, CL1, Para 4

Let Z be the set of all admissible policies; let Z_i be the subset of policies in Z that are integer-valued. Define the β -discounted, finite horizon, expected cost of policy z , when the system starts from state x at time $k = 0$, and is allowed to "move" for n steps (i.e., perform n transitions), as

(Eqn.(5))

where $L(z_k)$ is a linear function of the state trajectory and the control process z ; it has the interpretation of an instantaneous cost. A fairly general form for L , that fits our purposes is

(Eqn.(6))

where c, d are properly dimensioned vector constants. In resource allocation problems, where delay is the cost, we have $d = 0$; in pure blocking systems, we choose $c = 0$.

To show the exact dependence of $J_{\beta}(z, z)$ on x and z , let us rewrite (4) as

(Eqn.(7))

Then since x is constant and (Eqn.), where p denotes the probability distribution on Ω^n , we have

(Eqn.(8))

Equation (8) stresses the fact that the cost function is linear in the variables $z_k(w^h)$.^{*} The dependence of the cost on the probability distribution, the transitions and the constants c, d is "hidden" in $\gamma_k(w^h)$, to emphasize the dependence of the cost on the policy z . The exact form of $\gamma_k(w^h)$ can be routinely determined for the specific problem in hand [14]. We need only mention that $\gamma_k(w^h)$ is independent from the control policy z and the initial state x . For the purposes of the discussion in this section, the exact form of $\gamma_k(w^h)$ is irrelevant.

From (8) we see that the optimal policy is the one that minimizes the second term in the right hand side. From (7) the constraints fall in general into two categories:

(a) nonnegativity of states, namely

(Eqn.(9))

(b) boundedness of states, namely

(Eqn.(10))

where U is the bound. Since the constraints in (10) (\leq) are easily converted into constraints as in (9), we shall concentrate on constraints of the form (9) only.

Summarizing, the LP equivalent problem may take the form

$$\min eZ \quad (P)$$

$$AZ \leq b$$

This form is suitable to present results from sensitivity analysis.

Remark. The control variables are (Eqn.), and thus there is only a finite number of them. The constraint matrix A has elements that depend only on the transitions $\xi_k(w^h)$. The vector b depends only on the initial state z .

We have allowed $z_k(w^h)$ to take values in $[0,1]$. For sensitivity analysis, x , the initial state of the queueing system, should be also continuously-valued. In this case, the trajectory i will be continuously-valued; such a trajectory does not of course correspond to a real queueing system.

If, however, $x, z_k(w^h)$ are restricted to take integer-valued values only, then i will be integer-valued; in this case it does represent the evolution of the queueing system. The optimal cost function of the MDP in this case is given by^{*}

(Eqn.(11))

This is actually a problem in Integer Programming, the sensitivity analysis of which is not as well developed as that of a Linear Program. If we remove the restriction on integer-valued policies (and states), we have the above mentioned Linear Programming problem (P). Let

(Eqn.(12))

denote the optimal value function of problem (P). It is $W_n(z)$ for which results from sensitivity analysis apply. We wish to emphasize here that the functions W_n, V_n are quite different; first of all, they are even defined on different domains. If we can make, however, a suitable connection between them, then we can relate the properties of W_n (which we shall determine) to those of V_n (which we want).

Such a connection is indeed possible, if the Linear Program in (12) admits an integer-valued solution. In this case, for integer-valued x , (11) and (12) refer to the same problem. The optimal value function of the LP "contains" in some sense the optimal value of the MDP: we can recover $V_n(x)$ by "interpolating" $W_n(x)$ at the integer-valued points of its domain. Consequently, all the properties of $W_n(x)$ are automatically properties of $V_n(x)$ as well.

Vinioris: page 652, CL2, Para 8

Briefly, the procedure is as follows. From equation (1) (or (2)) the state is a linear function of the control actions z_k .

In reviewing the above, it can be seen that Vinioris teaches only the formulation of an MDP and the definition of a value function. However, the indicated locations in Vinioris cannot be interpreted as teaching the limitations of Applicant's claim directed "building approximations from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming" (which differs from the recitation of the limitation found in the Office Action).

In another example, the Office Action asserts that Schneider teaches "building a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming," at page 2726, CL1, Para 3 and 4. Further, the Office Action states that Schneider teaches "building one or more approximations from above and from below to a value function for the state using representations," at page 2722, CL1, Para 2 and page 2724, CL2, Para 6. Again, Applicant's attorney notes that the recitation of the limitations are incorrect. Moreover, at the indicated locations, Schneider merely states the following:

Schneider: page 2726, CL1, Para 3 and 4

Our experiments consider both deterministic and noisy versions of the problem. To build the deterministic version of the problem, we ran long (stochastic) simulations for each of the 421 actions and cached the mean observed production rate for each. For the noisy versions, we could have used noisy outcomes directly from the stochastic simulation, but instead we simply added Gaussian noise to the cached, deterministic production rates. This enabled our experiments to run significantly faster, and also allowed us to easily generate empirical results with varying amounts of noise.

Table 1 shows experimental results. The computation times reported are on a 200 MHz Pentium Pro. The first section contains results for the case where the factory output is deterministic and known. The purpose of the first two lines is to delimit the range of results we should expect from good algorithms. The "Random" algorithm builds a schedule by choosing 8 configurations at random, and it loses an enormous amount of money. Much of the cost is due to heuristic penalties for failing to satisfy customer demand.

Schneider: page 2722, CL1, Para 2

In this paper, we describe a Markov Decision Process (MDP) formulation of production scheduling which captures stochasticity, while retaining the ability to construct a schedule to meet demand forecasts. The solution to this MDP is a value function, specific to the current demand forecasts, which can be used to generate

optimal scheduling decisions online. We then describe an industrial application and a reinforcement learning method for generating an approximate value function in this domain. Our results demonstrate that in both deterministic and noisy scenarios, value function approximation is an effective technique.

Schneider: page 2724, CL2, Para 6

Here we describe a principled approach to generating closed-loop production scheduling policies with reinforcement learning methods. It combines the capabilities of both optimal control and AI search based methods. The approach is based on representing the problem as an MDP and representing the solution as an approximate value function. In contrast to many optimal control based methods, it produces a time-dependent policy specifically built to match current demand forecasts, rather than a time-invariant policy that ignores all demand information other than the current rate. Our experiments also demonstrate the ability to search hundreds of alternative factory configurations.

In reviewing the above, it can be seen that Schneider teaches only a Markov Decision Process (MDP) formulation of production scheduling which captures stochasticity, wherein the solution to the MDP is an approximate value function, specific to the current demand forecasts, which can be used to generate optimal scheduling decisions online. However, the indicated locations in Schneider cannot be interpreted as teaching "building approximations from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming" (which differs from the recitation of the limitation found in the Office Action).

Dangat and Hedlund fail to overcome these deficiencies in the combination of Viniotis and Schneider. Recall that Dangat and Hedlund were cited only against the dependent claims.

The various elements of Applicant's claimed invention together provide operational advantages over Viniotis, Schneider, Dangat, and Hedlund. In addition, Applicant's invention solves problems not recognized by Viniotis, Schneider, Dangat, or Hedlund.

Thus, Applicant submits that independent claims 1, 13, and 25 are allowable over Viniotis, Schneider, Dangat, and Hedlund. Further, dependent claims 2-12, 14-24, and 26-36 are submitted to be allowable over Viniotis, Schneider, Dangat, and Hedlund in the same manner, because they are dependent on independent claims 1, 13, and 25, respectively, and thus contain all the limitations of the independent claims. In addition, dependent claims 2-12, 14-24, and 26-36 recite additional novel elements not shown by Viniotis, Schneider, Dangat, or Hedlund.

III. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicant's undersigned attorney.

Respectfully submitted,

GATES & COOPER LLP
Attorneys for Applicant

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: June 11, 2004

By: 

Name: George H. Gates
Reg. No.: 33,500

GHG/